

Delineating Parameter Unidentifiabilities in Complex Models

Dhruva V. Raman, James Anderson, Antonis Papachristodoulou

October 28, 2016

Abstract

Scientists use mathematical modelling as a tool for understanding and predicting the properties of complex physical systems. In highly parameterised models there often exist relationships between parameters over which model predictions are identical, or nearly identical. These are known as structural or practical unidentifiabilities, respectively. They are hard to diagnose and make reliable parameter estimation from data impossible. They furthermore imply the existence of an underlying model simplification. We describe a scalable method for detecting unidentifiabilities, as well as the functional relations defining them, for generic models. This allows for model simplification, and appreciation of which parameters (or functions thereof) cannot be estimated from data. Our algorithm can identify features such as redundant mechanisms and fast timescale subsystems, as well as the regimes in parameter space over which such approximations are valid. We base our algorithm on a novel quantification of regional parametric sensitivity: multiscale sloppiness. Traditionally, the link between parametric sensitivity and the conditioning of the parameter estimation problem is made locally, through the Fisher Information Matrix. This is valid in the regime of infinitesimal measurement uncertainty. We demonstrate the duality between multiscale sloppiness and the geometry of confidence regions surrounding parameter estimates made where measurement uncertainty is non-negligible. Further theoretical relationships are provided linking multiscale sloppiness to the Likelihood-ratio test. From this, we show that a local sensitivity analysis (as typically done) is insufficient for determining the reliability of parameter estimation, even with simple (non)linear systems. Our algorithm can provide a tractable alternative. We finally apply our methods to a large-scale, benchmark Systems Biology model of NF- κ B, uncovering previously unknown unidentifiabilities.

1 Introduction

The increasing availability of computing power has motivated the mathematical modelling of complex systems in fields as diverse as systems biology, climate science, and economics. These models often incorporate many parameters, each representing an uncertain or varying quantity affecting model output. In such models it is often possible to change individual parameters by an arbitrarily large amount without affecting observed model output, so long as particular relationships between parameters are preserved. The involved parameters are then *structurally unidentifiable*, and cross-sections of parameter space preserving the aforementioned relationships are known as *structural unidentifiabilities*. Determination of such unidentifiabilities is an important, but hard problem in complex models. On the one hand it is a necessary prerequisite to parameter estimation from data: unidentifiable parameters cannot be estimated, and an attempt to do so may result in misleading information. On the other hand, knowledge of unidentifiabilities can provide mechanistic insight. For instance, suppose we know that the observed dynamics of a chemical reaction network model are affected by the product of two rate constants, but not their individual values. Then both rate constants are unidentifiable, while the associated unidentifiability is the product of the rate constants. Knowledge of the latter means that we can rewrite one parameter in terms of the other, both simplifying the model and informing us of a structural system property.

The determination of structurally unidentifiable parameters in nonlinear differential equation models has been an ongoing research topic for several decades. Multiple algebraic approaches, which directly analyse the system equations, have been suggested, e.g. [16, 29, 38, 23, 43, 26]. Several such approaches rely on the fact that the identifiability problem can be viewed as a special case of the observability problem [16], where parameters are considered as time-invariant states to be estimated. Research on observability in nonlinear systems often makes use of the Observability Rank Condition [16], which analyses the properties (specifically the rank) of a matrix of Lie derivatives of model output over time to determine unobservable model states [20, 22, 42]. Correspondingly, identifiability has been analysed using Differential Geometry-based approaches that exploit parametric symmetries inherent in the aforementioned matrix of Lie derivatives [38, 24, 13, 11, 12, 44]. Many of the approaches mentioned can potentially determine the exact functional form of the unidentifiabilities, but all suffer from a lack of general applicability and scalability, see e.g. [9]. The afore-

mentioned issue has motivated the implementation of numerical approaches based on model simulation at multiple points in parameter space, e.g. [2, 34]. Such approaches can be applied to much larger models, but cannot deal with general nonlinearities, and are not guaranteed to find all unidentifiabilities.

More recently, data-based approaches to identifiability analysis have been developed. The profile likelihood method [33] is a scalable numerical method that detects both unidentifiable parameters and associated unidentifiabilities, and has found widespread popularity in the Systems Biology community. It is also able to detect *practical unidentifiabilities* [39], in which large parameter perturbations induce small but nonzero changes in model output. However, it relies on iteratively moving individual parameters, while each time re-optimising model output over the other parameters. As such, a separate analysis is required for each parameter, and only one-dimensional unidentifiabilities are detected. The approach of [15] can flag the probable existence of more general unidentifiabilities through extensive model simulation over parameter space. However unidentifiabilities are provided only as probable functionally related groups of parameters; the functional relations themselves are not determined.

There is an underlying duality between the sensitivity of model predictions to parameter perturbation, and the uncertainty associated with parameter estimation in the presence of measurement uncertainty: it is hard to accurately estimate the model parameters from noisy data if a large change in some (combinations of) parameters induces only a small change in model output. Consequently, identifiability is closely related to the concept of model *sloppiness* [6, 14, 41, 37]. Here, the sensitivity of model predictions to infinitesimal parameter perturbation is highly anisotropic, in that the effect of perturbation in sensitive (‘stiff’) directions exceeds that of insensitive (‘sloppy’) directions by many orders of magnitude. It has been shown that sloppiness is a common feature in Systems Biology models [14]. A sloppy direction suggests the existence of a practical unidentifiability, although the precise correspondence has been debated [3, 10]. Conversely, if model sensitivity to parameter perturbation is dominated by a few ‘stiff’ directions in parameter space, it is possible that an underlying, macroscopic model simplification exists [36].

Existing literature [18, 40, 33, 17], has suggested that the local sensitivity characteristics of nonlinear models often poorly approximate their sensitivity to larger perturbations. This suggests that model sloppiness may not be informative in determining the uncertainty associated with parameter

estimates when measurement uncertainty is non-infinitesimal. Dually, it implies, when simplifying models, that the most appropriate reduced model is highly dependent on the range of parameters over which validity is intended. This case has been made in e.g. [1, 28]. We introduce a new notion, *multiscale sloppiness*, that quantifies sensitivity anisotropy as a function of the length scale of perturbation considered, relative to a fixed dataset/model prediction. We show how it furthermore relates to the geometry of the set of parameters satisfying a particular Likelihood-ratio hypothesis test. Multiscale sloppiness asymptotically corresponds to the standard formulation of model sloppiness in the limit of decreasing length-scale. We find that the sensitivity characteristics of models can alter drastically as perturbations of increasing magnitude are considered. In addition, both sloppiness and multiscale sloppiness can be highly dependent on the particular parameter vector considered, as we demonstrate subsequently by example. This suggests that caution must be exercised in labelling an entire model structure sloppy, based on analysis at a single parameter vector. Multiscale sloppiness allows for analysis of how the uncertainty region associated with a parameter estimate changes as the signal-to-noise ratio of the data decreases.

Our formulation of multiscale sloppiness leads to the presentation of a novel, numerical algorithm for unidentifiability detection. We take a particle in parameter space, and allow it to traverse parameter space via the solution of a set of Hamiltonian Equations. These equations are set such that the particle traces over structural unidentifiabilities if they exist, and otherwise over practical unidentifiabilities. Our algorithm has several attractive features. It is highly scalable, as shown in the examples. It does not require optimization of model output over parameter space, which can be computationally expensive and non-convex. Finally, it can detect not only the involved parameters, but also the *functional form* of the unidentifiability. The idea of analysing model structure through evolving particles in parameter space has been previously considered in [36], where particles evolve along geodesics of the Fisher Information Metric. These trajectories have different properties to ours: movement along such a curve changes the nominal parameter vector, while always considering infinitesimal magnitude parameter perturbations. Movement along our curves, by contrast, changes the length scale of parameter perturbation considered, relative to a fixed dataset or nominal parameter vector θ^* . More insight on this distinction is provided in Section 6.1. The method of [36] is, moreover, inapplicable when the model under analysis is structurally unidentifiable, and ill-conditioned when it has sloppy parameter vectors: calculation of the geodesic acceleration requires inversion

of the Fisher Information Matrix.

2 Quantifying Variability in Model Output over Parameter Space

We consider mathematical models incorporating a parameter vector θ , drawn from some set of allowable parameter vectors, Θ , known as the parameter space. Model output, as a function of the parameter vector, is denoted $y(\theta)$, and may be finite or infinite-dimensional. We often make use of a nominal (i.e. reference) output, which either corresponds to experimentally derived data D , or model output at a nominal parameter vector θ^* , i.e. $y(\theta^*)$. In order to quantify the degree of output disruption induced by parametric variation, we require a (scalar) cost function on parameter space. This is denoted $C_D(\theta)$ or $C_{\theta^*}(\theta)$, depending on the choice of reference. The lower the value of the cost function, the better that model output at θ represents the nominal output. The quantity $\nabla_{\theta}^2 C_{\theta^*}(\theta^*)$, which is used extensively throughout the paper, will be referred to as the *cost Hessian*. We note that the methods of our paper apply only to models for which both a cost function, and its gradient, can be calculated as a function of the parameters. However, our methods apply to any model for which these are calculable. We also assume that cost functions are twice differentiable.

We now describe a common, statistically motivated choice of cost function used in the parameter estimation problem. As previously stated, the methods of the paper are not restricted to this choice. If we consider measurement noise as a known random variable corrupting model output, then observations of model output are probabilistic. Given noise-corrupted data D , we can then take

$$C_D(\theta) = -\log L(\theta; D) + R(D), \quad (1)$$

where $L(\theta; D)$ represents the likelihood of θ , given D , and $R(D)$ is some offset term. Note that $L(\theta; D)$ is numerically equal to the probability density of the data D , given noise-corrupted model output at parameter vector θ . So parameter vectors that are more likely to produce the data D , when corrupted by measurement noise, have a lower cost. In fact, the parameter vector θ minimising $C_D(\theta)$ is known as the Maximum Likelihood Estimate (MLE).

If, instead of data D , we wish to explore model sensitivity around a nominal parameter vector θ^* , we can instead assume that hypothetical data $D(\theta^*)$ is generated as

$$D(\theta^*) = y(\theta^*) + \xi, \quad (2)$$

where ξ is a random variable denoting measurement noise. In this case we can take a cost function $C_{\theta^*}(\theta)$ as

$$C_{\theta^*}(\theta) = \mathbb{E}_{\theta^*}[-\log L(\theta; D(\theta^*))] - \mathbb{E}_{\theta^*}[\log L(\theta^*; D(\theta^*))]. \quad (3)$$

This represents the expected value of (1), assuming that data D was generated according to (2) and that $R(D) = -\mathbb{E}_{\theta^*}[\log L(\theta^*; D(\theta^*))]$. This choice of $C_{\theta^*}(\theta)$ is known as the Kullback-Leibler (KL) Divergence, or relative entropy, between the random variables $D(\theta^*)$ and $D(\theta)$, each generated according to (2). Moreover the cost Hessian $\nabla_{\theta}^2 C_{\theta^*}(\theta^*)$ then corresponds to the Fisher Information Matrix at θ^* .

Let us take a concrete example to provide context. Suppose that we assume that measurement noise ξ is Gaussian, being distributed as

$$\xi \sim \mathcal{N}(0, \Sigma),$$

for some covariance matrix Σ . Then the choices (1) and (3) of cost function are respectively given by

$$C_D(\theta) = \langle (D - y(\theta)), \Sigma^{-1}(D - y(\theta)) \rangle \quad (4a)$$

$$C_{\theta^*}(\theta) = \langle y(\theta^*) - y(\theta), \Sigma^{-1}(y(\theta^*) - y(\theta)) \rangle, \quad (4b)$$

where $\langle x, y \rangle$ denotes the inner product of vectors x and y . So this is the weighted sum of the squared residuals between the nominal output D or $y(\theta^*)$, and the model output $y(\theta)$.

Minimisers of the cost function represent model parameter vectors that ‘best approximate’ the nominal output. However, there are likely to be many other parameter vectors that are in reasonable agreement with the nominal output. We define

$$U_D(\epsilon) = \{\theta : C_D(\theta) \leq \epsilon\} \quad (5a)$$

$$U_{\theta^*}(\epsilon) = \{\theta : C_{\theta^*}(\theta) \leq \epsilon\}, \quad (5b)$$

as the ϵ -uncertainty regions, relative to D and θ^* respectively. A definition of (5), specialised to the case of Gaussian measurement noise as considered

in equations (4), was provided in [39]. When we take the costs (1) and (3), then the ϵ -uncertainty regions gain statistical properties, in both the Bayesian and Frequentist frameworks. We see this as sublevel sets of these cost functions precisely define the set of parameters whose likelihood given the (expected) data exceeds some critical value that is dependent upon the value of ϵ . In Bayesian statistics, an n -percent credible region in parameter space is defined as a set of parameters within which random samples of the posterior distribution would fall n times out of 100. An ϵ -uncertainty region of the form (5a) is then a *highest posterior density* credible region, under the assumption of a uniform prior. This type of credible region is a minimiser of the volume among the set of credible regions with the same percentage of credibility. In the presence of a more detailed prior distribution, the cost function (1) can be multiplied through by the prior for the highest posterior density property to hold. Meanwhile, the choice (3) of cost function represents those parameters whose expected likelihood, given data distributed according to (2), would place them within the highest posterior density credible region.

The Frequentist interpretation of ϵ -uncertainty regions is related to the task of hypothesis testing. Suppose, given data D , we took a null hypothesis of D being a noise-corrupted observation of $y(\theta^*)$ (and thus distributed according to (2)). Suppose further that we wanted to test the alternative hypothesis that D was generated from a different parameter vector within the parameter space. A Likelihood-ratio test (LRT) would reject the null hypothesis (i.e. θ^*) if the following condition was satisfied:

$$\Lambda_{\theta^*}(D) := \frac{L(\theta^*|D)}{\sup_{\theta \in \Theta} L(\theta|D)} \leq k^*, \quad (6)$$

for some critical value $k^* \in [0, 1]$, and where $\Lambda(D)$ is known as the LRT statistic. The significance level of this LRT is given as the α solving

$$Pr(\Lambda_{\theta^*}(D) \leq k^* | \theta^*) = \alpha. \quad (7)$$

Using the Likelihood-based cost function introduced in (4a), we can rewrite the LRT statistic as

$$\log(\Lambda_{\theta^*}(D)) = \inf_{\theta \in \Theta} C_D(\theta) - C_D(\theta^*).$$

Frequentist confidence regions are often generated by inverting the Likelihood-ratio test [8]. Specifically, one chooses the set of parameters that would have

passed the LRT given in (6), given the data, as the confidence region. This is given by

$$\begin{aligned} & \left\{ \hat{\theta} \in \Theta : C_D(\hat{\theta}) \leq \inf_{\theta \in \Theta} C_D(\theta) - \log(k^*) \right\} \\ &= U_D(\epsilon), \text{ for } \epsilon = \inf_{\theta \in \Theta} C_D(\theta) - \log(k^*). \end{aligned}$$

We see that $U_D(\epsilon)$ is precisely the confidence region gained by inverting a Likelihood ratio test. Moreover the data-independent cost function (3) can be written as $C_{\theta^*}(\hat{\theta}) = \mathbb{E}_{\theta^*}[C_D(\hat{\theta})]$, and we have:

$$\hat{\theta} \in U_{\theta^*}(\epsilon) \quad \Leftrightarrow \quad \mathbb{E}_{\theta^*}[\Lambda_{\hat{\theta}}(D)] \leq \exp(\epsilon),$$

where $\Lambda(D)$ is the Likelihood-ratio test statistic defined in (6). In other words, $U_{\theta^*}(\epsilon)$ precisely defines the set of parameter vectors that would be expected to pass a LRT, given noise-corrupted data generated from $y(\theta^*)$, and a significance level α related to ϵ by (7). An approximate probability distribution for $\Lambda_{\theta^*}(D)$ can be gained using Wilks' Theorem [8].

3 Relating Sloppiness and Unidentifiability

In this section we define model sloppiness and unidentifiability, and describe their relationship. Suppose two parameter vectors θ^* and θ have identical observed output. In this case, $C_{\theta^*}(\theta) := 0$, and we deem the model *structurally unidentifiable* at θ^* , following [4]. The particular parameters that differ between θ^* and θ are the structurally unidentifiable parameters. An example is the damped harmonic oscillator, modelled by the equation

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = 0. \tag{8}$$

Here $y(t)$ denotes vertical displacement as a function of time, and is taken as the observed variable. Meanwhile, $\theta = [m, c, k]$ denotes the mass, damping coefficient, and spring constant respectively, and dotted variables are time derivatives. Although there are three parameters, dividing (8) through by m shows that observations are invariant to parameter changes preserving the ratios $\frac{c}{m}$ and $\frac{k}{m}$. Thus, all three parameters are structurally unidentifiable, and we say that there is a structural unidentifiability over areas of parameter space preserving these ratios. As well as providing physical insight, this tells us that the parameter estimation problem is ill-posed for this model: we

cannot estimate θ by observing $y(t)$. A standard reparameterisation of (8) is

$$\ddot{y}(t) + 2\zeta\omega_0\dot{y}(t) + \omega_0^2 y(t) = 0. \quad (9)$$

Here ω_0 and ζ , known as the natural frequency and damping ratio respectively, are given by $\omega_0 = \sqrt{\frac{k}{m}}$, $\zeta = \frac{c}{2\sqrt{mk}}$. The reparameterised model is now identifiable, and the new parameters can be estimated from the observation data.

Even if a model is structurally identifiable at θ^* , which implies that $U_{\theta^*}(0) = \{\theta^*\}$, it may be that the set $U_{\theta^*}(\epsilon)$ remains large for small, but nonzero ϵ . This means that large tracts of parameter space induce very similar model output that cannot easily be distinguished from θ^* given even slightly noisy data. This is known as practical unidentifiability [39]: the effect on parameter estimation is that a small degree of measurement uncertainty results in a large degree of parametric uncertainty. Sloppy models often exhibit the same pathology ([14, 35]), and indeed there is a formal relationship between sloppiness and the properties of $U_{\theta^*}(\epsilon)$, which we present next.

Sloppiness quantifies the anisotropy in output disruption when infinitesimal parameter perturbations are applied to a nominal parameter vector θ^* . Specifically, the degree of sloppiness is the ratio between the degree of disruption in the most and least sensitive directions. Algebraically, this is the condition number of the cost Hessian $\nabla^2 C_{\theta^*}(\theta^*)$, i.e. the ratio of its maximal and minimal eigenvalues. Geometrically, this is the aspect ratio of the hyper-ellipse of perturbations $\delta\theta$ satisfying $\langle \delta\theta, [\nabla^2 C_{\theta^*}(\theta^*)] \delta\theta \rangle = 1$. The major and minor axes of this hyper-ellipse respectively represent the sloppiest and stiffest directions in parameter space. When the cost Hessian is singular, we will say that the degree of sloppiness is ‘infinite’. Note, however, that a singular cost Hessian does not imply local structural unidentifiability of a parameter vector, although the converse is true (see the model described in equation (14)). Furthermore if the cost Hessian is numerically generated, then it is not possible to discriminate between true singularity and the existence of very small but strictly positive eigenvalues.

Meanwhile the shape of $U_{\theta^*}(\epsilon)$, for locally structurally identifiable θ^* , is approximated by the same hyper-ellipse (up to scaling) in the limit of decreasing ϵ (see Appendix A and Figure 1). This is due to the effect of the second term of the Taylor expansion of $C_{\theta^*}(\theta)$ eventually dominating in (5b), as ϵ decreases. There is a direct statistical interpretation when the cost functions satisfy a condition known as asymptotic normality [8] (cost

functions satisfying this include (1), (3), (4a) and (4b)). This is the famous Cramer-Rao bound, which states that in the limit of decreasing measurement uncertainty, the matrix $[\nabla_{\theta}^2 C_{\theta^*}(\theta)]^{-1}$ and the covariance matrix of the MLE coincide. So sloppiness corresponds approximately to anisotropy in the MLE covariance. However this approximation can fare poorly when measurement uncertainty is non-negligible [18, 33].




Figure 1: As the length-scale ϵ tends to zero (dashed arrows), the shape of the ϵ -uncertainty region $U_{\theta^*}(\epsilon)$ is increasingly dictated by level sets of $\langle \delta\theta, \nabla_{\theta}^2 C_{\theta^*}(\theta^*) \delta\theta \rangle$ (dashed ellipse). This is demonstrated mathematically in Lemma 1 of Appendix A. The sloppiest and stiffest directions are the major and minor axes of the ellipse, respectively.

Sloppiness has been claimed to be a fundamental feature of model structure [41, 14], rather than a property of a particular parameter vector. Caution must be exercised in intuiting the former from the latter. Recall the oscillator model (9), and assume we observe the trajectory, while taking $\theta = [\zeta, \omega_0, y(0), \dot{y}(0)]$, so that initial position and velocity are now both parameters. We take a continuous-time analogue of the squared-residual cost function (4b):

$$C_{\theta^*}(\theta) = \int_0^\infty \|y(t, \theta) - y(t, \theta^*)\|_2^2 dt, \quad (10)$$

which can be reformulated as an algebraic function in the parameters using the methods of [31]. The degree of sloppiness at the two nominal parameter vectors $\theta^{1,*} = [0.5, 1, 1, 0]$; $\theta^{2,*} = [5, 1, 1, 0]$, are 94.0 and 3.41×10^7 respectively, a separation of six orders of magnitude.

4 Multiscale Sloppiness

We now introduce a new notion quantifying the anisotropy of model sensitivity for non-infinitesimal parameter perturbations. This furthermore leads towards an eventual algorithm for uncovering model unidentifiabilities. Define

$$D_{\bullet}^{max}(\delta) = \arg \max_{\theta} C_{\bullet}(\theta) : \|\theta - \theta^*\|_2^2 \leq \delta \quad (11a)$$

$$D_{\bullet}^{min}(\delta) = \arg \min_{\theta} C_{\bullet}(\theta) : \|\theta - \theta^*\|_2^2 = \delta, \quad (11b)$$

where $\|x\|_2$ denotes the Euclidean 2-norm of the vector x . Note that (11a) and (11b) are set-valued functions: the optimisers may not be unique. We can think of $D_{\bullet}^{max}(\delta)$ (resp. $D_{\bullet}^{min}(\delta)$) as the maximally (minimally) disruptive parameter vectors relative to θ^* at length scale δ . The dual, estimation-based interpretation follows naturally. Let us take $\bar{U}_{\bullet}(\epsilon)$ as the connected component of $U_{\bullet}(\epsilon)$ around θ^* . Given ϵ , we can take δ_1 and δ_2 such that $C_{\bullet}(D_{\theta^*}^{min}(\delta_1)) = \epsilon$, and $C_{\bullet}(D_{\theta^*}^{max}(\delta_2)) = \epsilon$. Then

$$D_{\bullet}^{min}(\delta_1) = \arg \max_{\theta} \|\theta - \theta^*\|_2^2 : \theta \in \bar{U}_{\bullet}(\epsilon) \quad (12a)$$

$$D_{\bullet}^{max}(\delta_2) = \arg \min_{\theta} \|\theta - \theta^*\|_2^2 : \theta \notin \bar{U}_{\bullet}(\epsilon). \quad (12b)$$

Thus they respectively signify the furthest parameter vectors inside, and the closest parameter vectors outside, the connected component of an ϵ -uncertainty region. Correspondingly, they represent the furthest (resp. closest) parameter vectors that would be expected to pass (fail) the LRT given in (6) (see Figure 2).

The previous discussion motivates our new definition of *multiscale sloppiness*, at length scale δ , relative to θ^* :

$$\mathcal{S}_{\bullet}(\delta) := \frac{\mathcal{S}_{\bullet}^{max}(\delta)}{\mathcal{S}_{\bullet}^{min}(\delta)} := \frac{\max_{\theta} C_{\bullet}(\theta) : \|\theta - \theta^*\|_2^2 \leq \delta}{\min_{\theta} C_{\bullet}(\theta) : \|\theta - \theta^*\|_2^2 = \delta}. \quad (13)$$

Note that $\lim_{\delta \rightarrow 0} \mathcal{S}_{\theta^*}(\delta)$ is the condition number of the cost Hessian: i.e. the traditional degree of sloppiness. There is no analogue for $\mathcal{S}_D(\delta)$, the data-dependent variant, as the traditional degree of sloppiness cannot account for experimental data. We see that $\mathcal{S}_{\theta^*}(\delta)$, for nonzero δ , quantifies anisotropy in the sensitivity characteristics of θ^* for parameter perturbations of nonzero




Figure 2: We pick balls of radius δ_1 and δ_2 (dashed) such that $\mathcal{S}_\bullet^{min}(\delta_1) = \mathcal{S}_\bullet^{max}(\delta_2) = \epsilon$. Square and circular dots respectively indicate locations of $D_\bullet^{max}(\delta_i)$ and $D_\bullet^{min}(\delta_i)$. The balls intersect the uncertainty region $U_\bullet(\epsilon)$ at $D_\bullet^{max}(\delta_1)$ and $D_\bullet^{min}(\delta_2)$.

length-scale. Meanwhile, the numerator and denominator of (13) are the costs associated with the maximally and minimally disruptive parameter vectors at length-scale δ . Their respective geometrical relationship with the ϵ -uncertainty region has been previously described, and is depicted in Figure 2.

Multiscale sloppiness can be highly dependent on δ . The left-hand side of Figure 3 shows the multiscale sloppiness of the damped oscillator model (9), with cost function (10), at one of the previously considered nominal parameter vectors. In this case, multiscale sloppiness increases with length scale δ . This is not a general rule, however. Consider the model

$$\begin{aligned} \begin{bmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \end{bmatrix} &= \begin{bmatrix} -1 & \phi_1 \\ \phi_1 + \phi_2 & \phi_2 - 4 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} & t \geq 0 \\ y(t) &= x_1(t) + x_2(t); & \theta = [\phi_1, \phi_2, x_1(0), x_2(0)] \end{aligned} \quad (14)$$

A graph of multiscale sloppiness against length scale for this model is given in Figure 3, for three nominal parameter vectors and the cost function (10). The first parameter vector is structurally identifiable (through the method of [4]), yet ‘infinitely’ sloppy, as $\nabla_\theta^2 C_{\theta^*}(\theta)$ is rank deficient. The second parameter vector is a small perturbation of the first, which has near identical multiscale sloppiness characteristics for nonzero δ , but has nonsingular $\nabla_\theta^2 C_{\theta^*}(\theta)$. The third parameter vector is much less sloppy at small length scales, but the situation reverses as δ grows. So if data were generated from

$\theta^{*,3}$ with enough measurement noise, the confidence region of the consequent parameter estimate would be expected to be more anisotropic than for $\theta^{*,1}$ and $\theta^{*,2}$, despite its lower level of sloppiness.




Figure 3: Graphs of $\mathcal{S}_{\theta^*}(\delta)$ against δ for (left) the damped oscillator (9) at $\theta^* = [0.5, 1, 1, 0]$, and (right) (14). The right-hand figure has three choices of nominal parameter vector: $\theta^{*,1} = [0, 0, 1, 2]$, $\theta^{*,2} = [0.001, 0, 1, 2]$, and $\theta^{*,3} = [-0.2, -0.2, 1, 2]$. Not shown on the right-hand graph is that for $\theta^{*,1}$, we have $\mathcal{S}_{\theta^*}(0) = \infty$.

Multiscale sloppiness is formulated with respect to the L_2 norm on parameter space (see equations (12a) and (12b)). Frequently, it is not easy to non-dimensionalise a model so that the parameters are dimensionless quantities. Often, the rescaling $\hat{\theta}_i = \log \theta_i$ is employed so that parametric variation is quantified as a relative, rather than absolute (see e.g. [14]), change, in this case. This method is inapplicable when a parameter component with a nominal value of 0 is considered, and may not be desirable in other circumstances. In such cases, the L_2 norm on parameter space depends on the units by which the parameters are measured, and can be weighted by rescaling. Thus the multiscale sloppiness, at a fixed δ , is sensitive to rescaling. However, analysis of the change in multiscale sloppiness properties as the length scale δ is varied yields model insights that are independent of the parameter units in a dimensionalised model. We now discuss the relationship between multiscale sloppiness, practical unidentifiability, and the scaling of units in parameter space.

If a modeller is interested in the behaviour of a model to perturbations in a parameter with units of distance, and these perturbations are on the length-scale of millimetres, then the modeller should use units of millimetres for the perturbation. If it turns out that only perturbations on the length-scale of metres have a discernible effect on model output, then the parameter is practically unidentifiable: it cannot be estimated to millimetre precision. Practical unidentifiability is a quantitative phenomenon, dependent on the requirements of the modeller, as opposed to structural unidentifiability, which is a structural phenomenon. In this case, minimally disruptive parameters will align with the co-ordinate axis of the parameter in question.

In general, practical unidentifiabilities cannot be explained as a consequence of poorly chosen units in parameter space, and point to more fundamental structural properties of the model. We provide several examples of such unidentifiabilities subsequently, in Subsection 6.3. We now provide a brief, motivating overview of one such example. Suppose a model contains a fast-timescale subsystem whose dynamics effectively equilibrate on the timescale of the observed dynamics. Increasing the (positive) time-constant τ of the subsystem will have almost no effect on dynamics. Decreasing the time-constant will also have initially negligible effect, until a critical value τ^* is reached at which instant equilibration starts to become a poor approximation. Past this point, perturbations to the time-constant will have a significant effect on observed model dynamics. Regardless of the units in which the time-constant is measured, we have one interval of τ values over which perturbation significantly affects dynamics, and another in which its effect is negligible. The minimally disruptive parameters for length-scale $\delta < \tau - \tau^*$, will be aligned with the co-ordinate axis of τ , and have negligible effect on the dynamics. The multiscale sloppiness will therefore be high. At $\delta \approx \tau - \tau^*$, the multiscale sloppiness will decrease, and the minimally disruptive parameters will diverge from the co-ordinate axis of τ , as further movement in the direction of decreasing τ becomes increasingly disruptive. Finally, a step change may occur, at which point the minimally disruptive parameters lie in a different direction of parameter space, reflecting a different model approximation.

5 Minimally Disruptive Curves in Parameter Space

In the previous section we highlighted the relationship between minimally disruptive parameter vectors and practically unidentifiable subspaces (as

quantified through the ϵ -uncertainty region). However, finding the minimally/maximally disruptive parameter vectors for a non-zero length-scale equates to a nonlinear, non-convex optimization which may be very difficult to solve. Therefore this section instead formulates the concept of minimally disruptive curves in parameter space, which can be traced by evolving a particle satisfying a set of Hamiltonian equations. These curves uncover regions of parameter space over which the cost function stays at, or close, to zero. Indeed these curves are guaranteed to trace over continuous branches of minimally disruptive parameters with strictly monotonically increasing length scale, where the latter exist. Hence, they are also guaranteed to trace over all structural unidentifiabilities with strictly monotonically increasing length-scale. A slight modification of the formulation is further guaranteed to trace over structural unidentifiabilities that do not satisfy the aforementioned strict monotonicity.

We first define the notion of a ‘continuous branch’ of minimally disruptive parameter vectors. Let us consider any continuous function Ξ_{min}^Δ satisfying:

$$\Xi_{min}^\Delta(\delta) \in D_{\theta^*}^{min}(\delta) \quad \forall \delta \in [0, \Delta], \quad (15)$$

where Δ is the largest considered length-scale. So such a Ξ_{min}^Δ exists when there is a continuous curve of minimally disruptive parameter vectors with strictly monotonically increasing distance δ from θ^* , up to distance Δ . We can pick an analogous function Ξ_{max}^Δ with the same properties, but formulated with respect to $D_{\theta^*}^{max}$ and the maximally disruptive parameters. Note that

$$\frac{d}{d\delta} [\Xi_{min}^\Delta(\delta)]_{\delta=0}, \text{ and } \frac{d}{d\delta} [\Xi_{max}^\Delta(\delta)]_{\delta=0}$$

lie in the eigenspace of the sloppiest and stiffest eigenvalues of the cost Hessian, respectively.

If a function $\Xi_{min}^\Delta(\delta)$ satisfying (15) exists, then it solves, for some F , the following optimization problem (see Appendix B):

$$\begin{aligned} \gamma^* &= \min_{\gamma \in \Gamma_{\theta^*}^F} : \int_{\gamma} C_{\bullet} : \\ \Gamma_{\theta^*}^F &= \left\{ \gamma : [0, F] \rightarrow \Theta : \begin{array}{l} \gamma(0) = \theta^*; \quad \frac{d}{ds} \|\gamma(s) - \theta^*\|_2 > 0 \\ \|\gamma'(s)\|_2 = 1 \end{array} \right\}. \end{aligned} \quad (16)$$

We therefore refer to a trajectory $\gamma(s)$ satisfying (16) as a *minimally disruptive curve* of length F , relative to θ^* . Note that structurally unidentifiable

parameter vectors are global minima of the cost function C_{θ^*} . So if there is a curve of structurally unidentifiable parameters in $\Gamma_{\theta^*}^F$, they will be minimally disruptive.




Figure 4: When a branch of minimally disruptive parameters $D_{\bullet}^{min}(\delta)$ vary continuously in δ , their union is a minimally disruptive curve. Each minimally disruptive parameter on the curve is the (non-unique) furthest point from θ^* within some (connected) ϵ -uncertainty region.

A naive approach to solution of (16) would be a gradient descent style algorithm, where $\gamma'(s)$ points along the direction of steepest descent of $C_{\bullet}(\gamma(s))$, subject to appropriate constraints. However in a structurally unidentifiable model, $C_{\theta^*}(\gamma(s))$ is always a local minimum of C_{θ^*} with null Jacobian $\nabla_{\theta} C_{\theta^*}$. Since exact nullity is never realised in a numerical context, a steepest descent direction will always exist and be highly sensitive to perturbation around the local minima, making numerical evolution ill-conditioned (see Figure 5). Instead we can consider (16) as a constrained variational problem (details in Appendix C), and obtain necessary conditions on $\gamma(s)$ to satisfy (16) by application of Pontryagin's Minimum Principle. These are listed below (and derived in Appendix C). Note that the derivative (in s) of a function f is denoted $f'(s)$, and the conditions must hold for all $s \in [0, F]$.

$$C_{\bullet}(\gamma(s)) + \langle \lambda(s), \gamma'(s) \rangle - C_{\bullet}(\gamma(F)) = 0 \quad (17a)$$

$$\mu_1(s)(\gamma(s) - \theta^*) + \left(C_{\bullet}(\gamma(F)) - C_{\bullet}(\gamma(s)) \right) \gamma'(s) = \lambda(s), \quad (17b)$$

where

$$\begin{aligned}\lambda(0) &= \left(C_{\bullet}(\gamma(F)) - C_{\bullet}(\gamma(0)) \right) \gamma'(0) \\ \lambda'(s) &= \mu_1(s) \gamma'(s) - \nabla_{\theta} C_{\bullet}(\theta) \quad \mu_1(s) \geq 0.\end{aligned}$$

Here $\lambda(s)$ is known as the costate, and is directly analogous to the momentum of Hamiltonian Mechanics. For a given parameter vector $\gamma(s)$, we assume that the cost function and its gradient, i.e. $C_{\bullet}(\gamma(s))$ and $\nabla_{\theta} C_{\bullet}(\gamma(s))$, can be calculated. If we additionally know $\lambda(s)$ and $C_{\bullet}(\gamma(F))$, then the remaining unknowns $\mu(s)$ and $\gamma'(s)$ are linear in (17), and can be solved for. Thus $\gamma(s)$ and $\lambda(s)$ uniquely evolve as a system of coupled differential equations in s , given a choice of $\gamma'(0)$ and $C_{\bullet}(\gamma(F))$. For (16) to hold as s approaches zero, $\gamma'(0)$ must be within the eigenspace of the smallest eigenvalue of the cost Hessian. Its magnitude is also fixed. Thus there are $2e$ choices of $\gamma'(0)$, where e is the dimension of the aforementioned eigenspace. Note that for $e > 1$, the $2e$ choices should form an orthonormal basis in this eigenspace, but such a basis is not unique. This situation occurs in the example of Subsection 6.2, which contains a two-dimensional structurally unidentifiable subspace. The discussion accompanying this example provides insight on the appropriate choice of basis.

The only free variable is now $C_{\bullet}(\gamma(F))$, which sets the initial momentum of the particle tracing the curve, and specifies the cost-cutoff at which the trajectory terminates. Note that this initial momentum ensures that we avoid the previously described pitfalls of a steepest descent approach to solving (16) (see Fig. 5).

A maximally disruptive curve can be defined by turning the minimisation of (16) into a maximisation. Such a curve maintains an analogous relationship with Ξ_{max}^{Δ} , if it exists. However it is of less interest from the point of view of unidentifiability characterisation. Note that in this case, a gradient-ascent style approach is sufficient from a computational viewpoint.

Note that model analysis using minimally disruptive curves is not restricted to models for which a continuous branch of minimally disruptive parameters (as given in (16)) exists. Instead, existence of the function given in (16) implies that the graph of minimally disruptive parameters is a minimally disruptive curve. An unidentifiability may not have strictly monotonically increasing distance from θ^* , making it untraceable by a minimally disruptive curve. In the case of structural unidentifiabilities, we can easily modify the optimization (16) to accommodate this. We would replace the condition




Figure 5: Bird’s eye view of level sets of the cost function. Darker areas have lower cost, while the black arrow traces the root of the cost function (i.e. a structural unidentifiability). Red lines denote the trajectory of a particle attempting to evolve numerically along the structural unidentifiability. On the left, a steepest descent style algorithm fails due to the stiffness of the direction of descent around a local minimum. On the right: the inherent momentum associated with trajectory evolution via solution of Hamiltonian equations ensures that numerical error does not completely disrupt evolution.

$\frac{d}{ds} \|\gamma(s) - \theta^*\|_2 > 0$ with

$$\frac{d}{ds} \|\gamma(s) - \theta^*\|_2 - c \|\gamma(s) - \theta^*\|_2 > 0, \quad (18)$$

for some constant $c \in [0, 2)$. This would trivially modify the Hamiltonian equations associated with the solution of (16), and derived in Appendix C. In this case, the optimization would be guaranteed to trace over any structurally unidentifiable curve $\gamma(s)$, whose angle $\phi(s)$ between $\gamma'(s)$ and the vector $\theta^* - \gamma(s)$ was always greater than $\cos^{-1}(\frac{c}{2})$. This guarantee is a result of global minimality of every structurally unidentifiable parameter vector, with respect to the cost function.

In the case of practical unidentifiabilities, the approach proposed in the previous paragraph is not guaranteed to be effective, although it pragmatically recovers highly practically unidentifiable subspaces not obeying monotonicity of distance from θ^* . For less sharply defined practically unidentifiable subspaces, usage of the condition (18) with $c > 0$ can result in a curve that confines itself to a ball around θ^* , outside of which all parameters have non-negligible cost. Thus, the curve will not be informative with respect to parameter perturbations with a larger length-scale than the ball’s radius. This can be heuristically tackled by increasing the initial momentum

imparted to the particle tracing the curve.

The formalism of minimally disruptive curves uncovers unidentifiabilities of a continuous nature. In this way, functional relations on parameter space over which model output changes little may be recovered. Models may also possess isolated parameter vectors that induce the same, or similar, output (see e.g. the example of Subsection 6.1). If so, the cost function will possess multiple isolated local minima. The ϵ -uncertainty regions may then consist of multiple, unconnected regions each surrounding a local minimum. Finding all local minima corresponds to a global optimization problem on the cost function, for which even evaluation often requires a numerical model simulation. This is in general an NP-hard problem [27]. While local minima of the cost function may well correspond to minimally disruptive parameters at the relevant length-scale, they are not recoverable through the generation of minimally disruptive curves.

6 Examples

We provide three examples of our methods. The first illustrates concepts introduced in previous sections, while highlighting the importance of non-locality in sensitivity quantification. The second verifies our methods on a benchmark model used in identifiability analysis. The third provides a more detailed analysis of a different benchmark example from the identifiability literature, specifically a metabolic reaction network. Novel structural and practical unidentifiabilities are identified. It is shown how our method can be used to gain additional mechanistic insight. Specifically, we uncover timescale-separated subsystems (and the regimes in which the timescale separation is valid), unnecessary mechanisms, and model approximations (together with their regimes of validity).

6.1 Illustrative Example

We first use a simple example to illustrate the concepts introduced in this paper, and compare them against the existing sloppiness framework. The model, previously considered in [35, 36], is given as:

$$y(\theta) = \left[e^{\frac{-\theta_1}{3}} + e^{\frac{-\theta_2}{3}}, e^{-\theta_1} + e^{-\theta_2}, e^{-3\theta_1} + e^{-3\theta_2} \right]. \quad (19)$$

$$C_{\theta^*}(\theta) = \|y(\theta) - y(\theta^*)\|_2^2$$

We take parameter space as $\Theta = [0, 4] \times [0, 4]$. We consider a nominal parameter vector $\theta^* = [4, 0.5]$. We deliberately place θ^* on the boundary of Θ so that only one direction of the sloppiest eigenvector of $\nabla_{\theta}^2 C_{\theta^*}(\theta)$ need be considered. This halves the analysis required without sacrificing illustrative power.

Note that model output (19) corresponds to observing a mixed exponential decay of the form $e^{-\theta_1 t} + e^{-\theta_2 t}$, at the timepoints $t = \{\frac{1}{3}, 1, 3\}$, and so model output is invariant with respect to permutation of the parameters. The model at θ^* is therefore structurally identifiable in a local neighbourhood but structurally unidentifiable from a global perspective (for nonzero parameter vectors). In such a case, a minimally disruptive curve is not guaranteed to uncover the structurally unidentifiable parameter vector $[0.5, 4]$: the guarantee only applies when structural unidentifiability is local. In this example, by chance, the minimally disruptive curve does in fact trace over $[0.5, 4]$.

We demonstrate in Figure 6 the differences inherent to evolving a particle in parameter space along its sloppiest eigenvector (which retains no memory of the initial vector θ^*), as compared to evolving a minimally disruptive curve (whose evolution is explicitly dependent on θ^*). To these ends, we first construct a vector field that, when evaluated at some $\theta \in \Theta$, points in the direction of the sloppiest eigenvector of $\nabla_{\theta}^2 C_{\theta}(\theta)$. The flow of this vector field, initialised from some θ^* (which we take to be $[4, 0.5]$), is compared to a minimally disruptive curve emanating from the same θ^* . We see that the former, although by definition always pointing in the sloppiest direction (to infinitesimal magnitude), diverges from the latter. Indeed, the flow of the vector field takes a higher cost route, as shown on the figure, tracing over parameter vectors with more highly diverging output (relative to $y(\theta^*)$) than the minimally disruptive curve (as required by definition). This discrepancy arises from the fact that the flow of the vector field, at some $\theta \in \Theta$ is determined independently of θ^* , and by analysis of the matrix $\nabla_{\theta}^2 C_{\theta}(\theta)$. The direction minimising local change in $\nabla_{\theta}^2 C_{\theta}(\theta)$, is not necessarily the direction minimising non-local change in $C_{\theta^*}(\theta)$. Furthermore the minimal eigenvector of the cost Hessian at the nominal, i.e. $\nabla_{\theta^*}^2 C_{\theta}(\theta)$, is a poor marker of the direction of minimal sensitivity to non-infinitesimal parameter perturbation.




Figure 6: Left: Vector field formed by sloppiest eigenvectors of the cost Hessian (one direction only), for the model (19). Overlaid are (solid line, red) a curve of minimally disruptive parameters, and (broken line, purple) the flow of the vector field. Right: The costs associated with these two curves. The x-axis charts Euclidean distance from θ^* , while the y-axis charts the associated cost.

6.2 IL13-induced JAK-STAT Pathway

We next verify our methods on a differential equation model of the IL13-induced JAK-STAT pathway [30], which was used as a benchmark for identifiability analysis in [32]. The model has a 23-dimensional parameter space (further details in the Supplementary Information) together with published estimates of each parameter, constituting θ^* . In [32], two structural unidentifiabilities were found: model output was invariant over sets of the form

$$\theta_{17}^* \theta_{22}^* = \theta_{17} \theta_{22}; \quad (20a)$$

$$\theta_{15}^* \theta_{21}^* = \theta_{15} \theta_{21} \text{ and } \theta_{15}^* \theta_{11} = \theta_{11}^* \theta_{15}. \quad (20b)$$

Only the method of [33] succeeded in detecting both the parameters involved in (20a), (20b), and their functional relations. This method requires evolution of a separate curve for each of the 23 parameters. To generate the i^{th} curve, θ_i^* is varied incrementally, and the other parameters iteratively re-optimised so as to minimise the effect on the cost function. This has the advantage of not requiring computation of $\nabla_{\theta} C_{\bullet}$ (unlike our method), and also generates componentwise confidence intervals. However the re-optimization steps can be expensive.

Our method recovered both (20a) and (20b), requiring only four curve generations (evolvable in parallel). The Cost Hessian, at θ^* , had two eigenvectors with numerically zero eigenvalues. Therefore the four curves had their initial velocity $\gamma'(0)$ aligned in both the positive and negative directions of each of these eigenvectors. The trajectories are shown in Figure 7. The top half of Figure 7a shows the log-space evolution of each component of $\gamma(s)$, relative to θ^* . The components corresponding to the rate constants θ_{17} and θ_{22} are marked, and we see that the summation $\log \theta_{17} + \log \theta_{22}$ is preserved to numerical precision. The lower half of the figure shows, meanwhile, that the cost $C_{\theta^*}(\gamma(s))$ is zero, to numerical precision. This implies a structural unidentifiability over areas of parameter space preserving the product $\theta_{17}\theta_{22}$, as given in (20a). Similarly, the trajectory in Figure 7b preserved the expressions $\log \theta_{15} - \log \theta_{11}$, and $\log \theta_{15} + \log \theta_{21}$, implying the unidentifiability of (20b).

The fact that the JAK-STAT model has two, one-dimensional, structural unidentifiabilities implies that the Cost Hessian has a two-dimensional null-space. As stated, two eigenvectors spanning the null-space set the initial directions $\gamma'(0)$ of the four generated curves. It is possible, however, to synthesise multiple such eigenvector pairs, as a two dimensional linear space does not have a unique orthonormal basis. Using the premise that different structural unidentifiabilities will involve different parameters, we used an eigenvector pair whose nonzero entries did not correspond at all, resulting in the curves of Figure 7 and the unidentifiabilities of (20a) and (20b). If we denote the eigenvectors chosen as v_1 and v_2 , then the choice $\tilde{v}_1 = v_1 + \frac{2}{3}v_2$ and $\tilde{v}_2 = v_1 - v_2$ also forms an orthonormal basis. Both, however, have nonzero components for all five parameters involved in the structural unidentifiabilities. Generating minimally disruptive curves whose initial directions aligned with \tilde{v}_1 and \tilde{v}_2 yielded different structurally unidentifiable curves, which involved all five parameters, but satisfied (20a) and (20b). The latter relations could be uniquely obtained from the curves by applying basic linear algebra to the linear relations preserved in the logarithms of the parameters.

6.3 NF- κ B Regulatory Module

We next find novel structural and practical unidentifiabilities in a benchmark differential equation model used for identifiability analysis in the literature: the 29-parameter model of the NF- κ B regulatory module [21]. The model describes an intracellular mechanism mediating dynamics of the NF- κ B pro-

tein complex. A schematic diagram, based on one provided in [21], is given in Figure 8, and details on the system equations and parameters are provided in the Supplementary Information. The system is excited by extracellular TNF. Previous analyses [9, 21] declared the model structurally identifiable, under relaxed assumptions necessary to maintain viability of the respective algorithms. Specifically, identifiability held so long as all parameters could be discriminated on the basis of some (unknown) time-course of TNF excitation, rather than the binary TNF signal assumed in [21]. In [9], sixteen parameters were furthermore fixed to maintain computational tractability. We analyse the full model, taking θ^* as the (complete) set of published parameter estimates/assumptions provided in [21].

We evolved two curves satisfying (17) in parallel, setting $\gamma'(0)$ in both the positive and negative directions of the eigenvector corresponding to the minimal eigenvalue of the cost Hessian. The result is shown in Figure 9a, which connects both curves at θ^* . The top half of the figure shows the log-space evolution of each component of $\gamma(s)$, relative to θ^* . The components corresponding to the rate constants k_2 and c_4 are marked, and we see that the summation $\log k_2 + \log c_4$ is preserved to numerical precision. The lower half of the figure shows, meanwhile, that the cost $C_{\theta^*}(\gamma(s))$ is zero, to numerical precision. This implies a structural unidentifiability over areas of parameter space preserving the product $k_2 c_4$. This makes sense, as k_2 and c_4 respectively represent the translation (i.e. production) rate of the protein A20, and the inactivation rate of the kinase IKKa caused by A20. Since A20 and its transcript are not directly observed, it is clear that A20 only enters the model through its effect on IKKa. This refines the conclusion of [21], where the effect of A20 on dynamics was proposed to be a combination of its *concentration* (rather than translation rate), and the inactivation rate.

We next rewrote k_2 as $k_2 = \frac{k_2^* c_4^*}{c_4}$, thus reducing the dimensionality of parameter space by one, and eliminating the previous structural unidentifiability. Two minimally disruptive curves were again evolved in parallel, with initial directions in the positive and negative directions of the sloppiest eigenvector. Results are shown in Figure 9b. Three parameters show significant movement, and are involved in a practical unidentifiability. They are listed below:

- $e2a$ represents the rate of transport of the $I\kappa B\alpha$ —NF- κ B complex from the cytoplasm to the nucleus. This was appreciated as a fast-timescale reaction in [21], as shown in Figure 8.

- c_{5a} is the constitutive degradation rate of free cytoplasmic $I\kappa B\alpha$.
- t_1 is the rate constant for catalysis of the $I\kappa B\alpha$ -IKK α complex into IKK α , also described as a fast timescale reaction in [21].

By varying these three parameters in turn, we see that each individual parameter has a close to zero effect on model output when varied, and the unidentifiability does not depend on their correlated change. Since e_{2a} and t_1 are fast-timescale reaction rates, the involved reactants and products are equilibrated near-instantaneously at θ^* , and increasing the rates makes no difference to observed model dynamics. Similarly, Figure 8 implies that $I\kappa B\alpha$ -IKK α degradation is dominated by terms other than constitutive degradation, and decreasing c_{5a} will therefore barely affect observed dynamics. Thus, elimination of the reaction corresponding to the rate c_{5a} , and instant equilibration of the reactions corresponding to t_1 and c_{5a} , would negligibly affect dynamics at θ^* .

Note from Figure 8 that as the minimally disruptive curve travels in the direction of decreasing e_{2a} , decreasing t_1 , and increasing c_{5a} , the associated cost starts to increase, until a step change in behaviour occurs. This can be explained mechanistically. Decrease in t_1 and e_{2a} slows the timescale of their associated reactions, until instant equilibration is no longer a valid approximation and observed dynamics are affected. Meanwhile, as c_{5a} increases it starts to make a meaningful contribution to $I\kappa B\alpha$ degradation, and thereby also affects observed dynamics. Crucially, our algorithm finds at what parameter values these step changes occur, thereby defining the areas of parameter space over which the instant equilibration/zero degradation approximations are valid.

We fixed the parameters c_{5a} , e_{2a} , and t_1 and again evolved minimally disruptive curves as previously. Results are shown in Figure 9c. We see that in a region around θ^* , dynamics are insensitive to parameter changes preserving the product $c_{1a}c_{4a}$. The former is the rate constant for promotion of the mRNA transcript for $I\kappa B\alpha$ by free nuclear NF- κ B. The latter is the translation rate for $I\kappa B\alpha$. Since the production rate of a protein is the product of its transcription and translation rates, then, as long as mRNA transcription is dominated by c_{1a} , it makes sense that preservation of the product $c_{1a}c_{4a}$ should not affect observed dynamics. However, the figure shows that this unidentifiability breaks down when c_{1a} decreases sufficiently. At this point, the cost starts to increase, and a step change in the practical unidentifiability is observed. The reason is that the model incorporates constitutive transcription of $I\kappa B\alpha$, which becomes the dominant contributor to overall

transcription in this regime.

Again, we removed the practical unidentifiability in $c_{1a}c_{4a}$ by substituting $c_{1a} = \frac{c_{1a}^*c_{4a}^*}{c_{4a}}$, and evolved minimally disruptive curves (Figure 9d). The resulting curve involved a_2 , which represents the association rate of IKKa and $I\kappa B\alpha$ in the nucleus. In [21], it is depicted as a fast timescale reaction, with a nominal value of 0.2 obtained by assumption. In fact, as shown in Figure 9d, we can decrease a_2 by over twenty orders of magnitude without affecting nominal dynamics, so long as other parameters are changed in a compensatory manner. Thus we could switch off this mechanism with only a negligible change in nominal dynamics.

For our final curve evolution, we fixed a_2 , with results shown in Figure 9e. This results in a complicated practical unidentifiability involving many parameters changing by several orders of magnitude without perceptibly affecting dynamics. Iteratively fixing three further parameter vectors whose values changed considerably in Figure 9e also resulted in complicated practical unidentifiabilities involving large parameter variations with little effect on the cost (these are not shown graphically). We could not attach mechanistic meaning to the practical unidentifiability of Figure 9e, and indeed it may not exist (see e.g. [5] for a more detailed discussion).

To demonstrate, suppose that we only measured model output at a small subset of timepoints. A whole host of new unidentifiabilities would arise, whose particular form would be dependent on the subset chosen, but which would not reflect mechanistic features of the model. These would be eliminated by measurement at more timepoints. Similarly, extra information from the NF- κ B model may be required to sufficiently constrain its parameter space. For instance, a knockdown experiment inhibiting expression of a particular gene could be conducted, and the appropriately modified model fitted to the consequent data. We could then require that for a parameter vector to be low-cost, it must not only recreate a set of nominal dynamics, but also behave appropriately when modified to model the knockdown experiment.

The modeller may wish to estimate a particular subset of the parameter vector. In this case, further generation of minimally disruptive curves, and consequent fixation of parameters, would be recommended. At some point, minimally disruptive curves would start to induce significant perturbations of the cost function over small length-scales of perturbation. At this point a set of parameters would have been found whose *a priori* fixation would

result in well-posed parameter estimation for the remaining parameters of the model.

7 Discussion

We have provided a numerical approach to elucidation of both unidentifiable parameters and their associated unidentifiabilities. It can complement symbolic approaches to finding unidentifiable parameters, which additionally provide theoretical guarantees, but can suffer from scalability issues. In particular, algebraic approaches based on the Observability Rank Condition [16] are in general computationally intensive even for modestly-sized models, but work well when there are few parameters and many outputs under consideration. Our method can be used as a preconditioning tool for such algorithms, as it can highlight groups of parameters likely to be structurally unidentifiable. All other parameters can then be fixed, and unidentifiability of the simplified model checked algebraically. For instance, in the example of the NF- κ B network from [21], we fixed all parameters except k_2 and c_4 (which we discovered were structurally unidentifiable through our algorithm, as discussed previously and shown in Figure 9a). The reduced, two-parameter model could then be easily analysed algebraically, and verified as structurally unidentifiable by direct application of the Observability Rank Condition. The same, integrated procedure was also performed successfully to verify the structurally unidentifiable parameters of the JAK-STAT model analysed previously.

Our method additionally inferred the functional form of the structural unidentifiabilities, through analysis of the trajectories generated by the algorithm. In general, this inference may not be trivial. It is considerably simplified by considering the trajectories in log-space, as depicted in Figure 9, so long as the unidentifiability is rational. The reason is that if there is an unidentifiability over a relation of the form $p(\theta) = 0$, where $p(\theta)$ is rational, then a linear function of the logarithms of the parameters is preserved, and inference of the functional relation becomes a linear regression problem.

A key theme of this paper is that the parametric sensitivity characteristics of a mathematical model can depend highly on both the parameter vector considered and the length-scale of parameter perturbation. We saw in the NF- κ B network example that there were several practical unidentifiabilities that held in a certain region of parameter space, but had a distinct boundary

at which a step change in the unidentifiability occurred (see e.g. Figure 9b, Figure 9c). These corresponded to the regions in which particular model approximations were valid. For instance, some such unidentifiabilities involved the rate constants of fast-timescale reactions, which were unidentifiable as long as a timescale separation existed between the reaction dynamics and observed model dynamics, but regained identifiability once they fell below a certain level. Our algorithm identified both the practical unidentifiabilities, and the boundaries in parameter space at which the form of the unidentifiabilities experienced a qualitative change. It is therefore of use in the model reduction problem, both for identifying redundant mechanisms, and quantifying the regions of parameter space at which they remain redundant.

The majority of the computational time is taken up simulating the model at different points in parameter space, in order to calculate the associated cost function, and gradient. The density of simulations on a minimally disruptive curve varies with the ‘momentum’ (i.e. the choice of magnitude of the initial costate $\lambda(0)$), the choice of ODE solver used to evolve the Hamiltonian flow (we used the ‘ode113’ solver from the MATLAB toolbox [25]), and the model in question. We took $C_{\theta^*}(\gamma(F)) = 0.1$ to set our momentum, for all curves generated in the previous example. 440 internal evaluations of the cost function and its gradient were required by the ode113 solver to produce Figure 9a, while Figure 9e required 1363 evaluations. Each combined evaluation of the cost function and gradient took about five seconds on a single-core, 3.40 GHz processor. The gradient was calculated by solving the sensitivity equations associated with the model (see [19]), rather than by finite-difference methods.

Consideration of model sloppiness at a non-local level has been considered previously for locally identifiable models (see [37, 35]). These approaches come from a different angle, and consideration of the varying purposes behind our formulation and theirs is warranted. In the approach of [37], the cost Hessian is taken as a Riemannian Metric on parameter space. The distance between parameter vectors is then taken as the length of a geodesic with respect to this metric. Given some nominal parameter vector θ^* , geodesics with initial direction along the eigenvalues of the cost Hessian are evolved until the borders of parameter space are reached. The ratio of the longest and shortest geodesic lengths then determines the global sloppiness of the model.

The global sloppiness quantification outlined has the advantage of being independent of the parameter units: geodesic lengths are invariant to the

choice of co-ordinates. They do, however, depend on the choice of parameter θ^* from which they are initialised, which highly affects the anisotropy of the cost Hessian, as previously demonstrated for the oscillator model given in (9). Evolution of the geodesic relies on successive linearisations of the cost function relative to the current location of the geodesic, rather than θ^* , whose dynamics are ‘forgotten’. Thus it is determined by the infinitesimal sensitivity properties of the model, relative to a constantly changing nominal parameterisation. We instead consider non-infinitesimal sensitivity properties relative to a fixed nominal parameterisation, or data, which allows for a link between multiscale sloppiness and the confidence regions of parameter estimates fitted to significantly noisy data. The distinction was discussed in Subsection 6.1, using a simple exponential decay model taken from [37] (see also Figure 6). Note that minimally disruptive curves and geodesics of the cost Hessian are relevant on different classes of model. Minimally disruptive curves are formulated specifically to recover structural and practical unidentifiabilities. Calculation of the acceleration term for the geodesic described, by contrast, involves inversion of the cost Hessian. This is by definition singular in locally structurally unidentifiable models, and numerically singular in locally practically unidentifiable models, and so cannot be inverted.

8 Conclusion

In conclusion, we have presented a scalable algorithm allowing for systematic characterisation of the structural and practical unidentifiabilities in a model. The approach is conceptually novel, treating unidentifiable cross-sections of parameter space as solutions of an optimal control problem, which is then solved numerically. En route, we have formulated a new approach to quantifying model sensitivity to parameter perturbations of non-infinitesimal magnitude. In the dual regime of estimation, we showed how this corresponds to the confidence regions of parameter estimates in the presence of non-infinitesimal measurement noise. The scalability, accuracy, and utility of the approach have been demonstrated through the discovery of new unidentifiabilities in benchmark models. Our hope is that once a model of a process is constructed, our methods can be used to systematically identify hidden unidentifiabilities and mechanistic redundancies. Subsequent parameter estimation and model-based analysis of the process is then made more tractable and reliable.

A The correspondence between structural unidentifiability, sloppiness, and multiscale sloppiness

In the main text we define a non-negative cost function $C_{\theta^*} \in \mathcal{C}^2$ on parameter space, such that $C_{\theta^*}(\theta^*) = 0$. Here \mathcal{C}^k denotes the space of k -times differentiable, continuous functions. We then define multiscale sloppiness, for parameter perturbations of length scale δ , and in the absence of data, as

$$\mathcal{S}_{\theta^*}(\delta) := \frac{\mathcal{S}_{\theta^*}^{max}(\delta)}{\mathcal{S}_{\theta^*}^{min}(\delta)} := \frac{\max_{\theta} C_{\theta^*}(\theta) : \|\theta - \theta^*\|_2^2 \leq \delta}{\min_{\theta} C_{\theta^*}(\theta) : \|\theta - \theta^*\|_2^2 = \delta}$$

$$\mathcal{S}_{\theta^*}(0) := \lim_{\delta \rightarrow 0} \mathcal{S}_{\theta^*}(\delta).$$

We state that $\mathcal{S}_{\theta^*}(0)$ agrees with the traditional quantification of sloppiness at θ^* , i.e. the ratio of the maximum to the minimum eigenvalue of $\nabla^2 C_{\theta^*}(\theta^*)$. We now prove the claim. Since θ^* is a local minimum of C_{θ^*} , we have both that $\nabla_{\theta} C_{\theta^*}(\theta^*) = 0$, and that $\nabla_{\theta}^2 C_{\theta^*}(\theta^*)$ is positive semidefinite. If we consider a perturbation $\delta\theta$, the remainder theorem ensures

$$C_{\theta^*}(\theta^* + \delta\theta) = \langle \delta\theta, \nabla^2 C_{\theta^*}(\theta^*) \delta\theta \rangle + O(\|\delta\theta\|_2^3). \quad (21)$$

Therefore, for $\|\delta\theta\|_2^2 \leq \delta$, we have

$$\lim_{\delta \rightarrow 0} \frac{\mathcal{S}_{\theta^*}^n(\delta)}{\mathcal{S}_{\theta^*}^d(\delta)} = \lim_{\delta \rightarrow 0} \frac{\max_{\delta\theta} \langle \delta\theta, \nabla^2 C_{\theta^*}(\theta^*) \delta\theta \rangle}{\min_{\delta\theta} \langle \delta\theta, \nabla^2 C_{\theta^*}(\theta^*) \delta\theta \rangle} : \quad (22)$$

Clearly the numerator and denominator of (22) must respectively lie in the eigenspace of the maximal and minimal eigenvectors of $\nabla^2 C_{\theta^*}(\theta^*)$.

We state in the main text that the shape of the ϵ -uncertainty region $U_{\theta^*}(\epsilon) := \{\theta : C_{\theta^*}(\theta) \leq \epsilon\}$, approximates a hyper-ellipse in the limit of decreasing ϵ for structurally identifiable models. Let us formalise this statement mathematically.

Lemma 1. *Limiting behaviour of $U_{\theta^*}(\epsilon)$* *Given a structurally identifiable parameter θ^* , and a compact parameter space Θ , consider the volume-normalised version of $U_{\theta^*}(\epsilon)$ given by*

$$\tilde{U}_{\theta^*}(\epsilon) = \left\{ \theta^* + \frac{\delta\theta}{\sqrt{2\epsilon}} : \theta^* + \delta\theta \in U_{\theta^*}(\epsilon) \right\}. \quad (23)$$

Then

$$\lim_{\epsilon \rightarrow 0} \tilde{U}_{\theta^*}(\epsilon) = \theta^* + \left\{ \delta\theta : \langle \delta\theta, \nabla^2 C_{\theta^*}(\theta^*) \delta\theta \rangle \leq 1 \right\}.$$

Proof. Since θ^* is structurally identifiable, we know that $U_{\theta^*}(0) = \{\theta^*\}$. Our first step is to show, for sufficiently small ϵ , that $U_{\theta^*}(\epsilon)$ is a connected set. If not, then there exists a sequence $\{\theta^{(i)}\}_{i=1}^\infty$ such that, for all $i \in \mathbb{N}$:

$$\exists \kappa > 0 : \|\theta^{(i)} - \theta^*\|_2^2 > \kappa \text{ and } C_{\theta^*}(\theta^{(i)}) < \frac{1}{i}.$$

The Extreme Value Theorem then guarantees the existence of a limiting parameter vector $\theta^{(\infty)} \in \Theta$ satisfying

$$\exists \kappa > 0 : \|\theta^{(\infty)} - \theta^*\|_2^2 > \kappa \text{ and } C_{\theta^*}(\theta^{(\infty)}) = 0,$$

contradicting structural identifiability.

Since $U_{\theta^*}(\epsilon)$ is connected for sufficiently small ϵ , and C_{θ^*} is continuous, we know that

$$\lim_{\epsilon \rightarrow 0} \left[\max_{\delta\theta \in \Theta} \|\delta\theta\|_2^2 : C_{\theta^*}(\theta^* + \delta\theta) < \epsilon \right] = 0.$$

Next, taking a Taylor expansion of (5b), and using the definition of \tilde{U}_{θ^*} from (23), we see that

$$\tilde{U}_{\theta^*}(\epsilon) = \left\{ \delta\theta : \left(\frac{(\sqrt{2\epsilon})^2}{2} \langle \delta\theta, [\nabla_{\theta}^2 C_{\theta^*}(\theta^*)] \delta\theta \rangle \right) + \mathcal{O}(\|2\sqrt{\epsilon}\delta\theta\|_2^3) \leq \epsilon \right\}.$$

We can divide the above equation through by ϵ , and note, by definition, that

$$\lim_{\epsilon \rightarrow 0} \frac{\mathcal{O}(\|2\sqrt{\epsilon}\delta\theta\|_2^3)}{\epsilon} = 0,$$

to get

$$\lim_{\epsilon \rightarrow 0} \tilde{U}_{\theta^*}(\epsilon) = \{ \delta\theta : (\langle \delta\theta, [\nabla_{\theta}^2 C_{\theta^*}(\theta^*)] \delta\theta \rangle) \leq 1 \}$$

as required. \square

The main text states that structural unidentifiability is equivalent to the existence of a strictly positive δ such that $\mathcal{S}_{\theta^*}(\delta)$ is infinite. We prove the forward implication first. Structural unidentifiability implies there exists $\hat{\theta} \neq \theta^*$ such that $C_{\theta^*}(\hat{\theta}) = 0$. So if we take $\delta = \|\hat{\theta} - \theta^*\|_2^2$, then we are assured that $\delta > 0$. But this means that $\mathcal{S}_{\theta^*}^d(\delta) = 0$. We now prove the backward implication. Continuity of C_{θ^*} is assumed. Therefore C_{θ^*} must

attain a finite maximum value on any compact set, by the Extreme Value Theorem. This implies that $\mathcal{S}_{\theta^*}^n(\delta)$ is finite. So $\mathcal{S}_{\theta^*}(\delta)$ can only be infinite, if $\mathcal{S}_{\theta^*}^d(\delta) = 0$. Therefore there exists a $\hat{\theta}$ such that $\|\hat{\theta} - \theta\|_2^2 = \delta$ and $C_{\theta^*}(\hat{\theta}) = 0$. Since $\delta > 0$, we are assured that $\hat{\theta} \neq \theta^*$, and thus $\hat{\theta}$ satisfies the conditions for structural unidentifiability.

B The correspondence between minimally disruptive parameters and minimally disruptive curves

The Lemma below shows that continuous branches of minimally disruptive parameters are minimally disruptive curves.

Lemma 2. *Suppose a function $\Xi_{min}^\Delta(\delta)$ satisfies*

$$\Xi_{min}^\Delta(\delta) \in D_{\theta^*}^{min}(\delta) \quad \forall \delta \in [0, \Delta], \quad (24)$$

for some $\Delta > 0$, and where $D_{\theta^}^{min}(\delta)$ denotes the minimally disruptive parameter at length-scale δ . Then $\Xi_{min}^\Delta(\delta)$ is a solution of the following optimization problem:*

$$\begin{aligned} \gamma^* = \min_{\gamma \in \Gamma_{\theta^*}^F} : \int_{\gamma} C_{\bullet} : \\ \Gamma_{\theta^*}^F = \left\{ \gamma : [0, F] \rightarrow \Theta : \begin{array}{l} \gamma(0) = \theta^*; \\ \|\gamma'(s)\|_2 = 1 \end{array} \quad \frac{d}{ds} \|\gamma(s) - \theta^*\|_2 > 0 \right\}. \end{aligned} \quad (25)$$

Proof. First note that for any $\gamma \in \Gamma_{\theta^*}^F$, we can take the reparameterisation

$$\delta(s) = \|\gamma(s) - \theta^*\|_2. \quad (26)$$

We have, for $s > 0$, that $\frac{d}{ds} \|\gamma(s) - \theta^*\|_2 > 0$. This ensures that $\delta(s)$ is a bijective function, and so it has an inverse. We can therefore parameterise the line integral of (25) using either δ or s . Note however that the reparameterisation $\delta(s)$ will depend on the specific curve γ .

Suppose that a function Ξ_{min}^Δ satisfying (24) exists, but is not a minimiser of (25). We take the actual optimiser of (25) as γ^* . Reparameterisations of the form (26) will be different for Ξ_{min}^Δ and γ^* , so we denote them $\delta^1(s)$ and $\delta^2(s)$ respectively. Now by assumption, we have that

$$\int_0^{\delta^1(F)} C_{\bullet}(\gamma^*(\delta)) \left\| \frac{d}{d\delta} [\gamma^*(\delta)] \right\|_2 d\delta < \int_0^{\delta^2(F)} C_{\bullet}(\Xi_{min}^\Delta(\delta)) \left\| \frac{d}{d\delta} [\Xi_{min}^\Delta(\delta)] \right\|_2 d\delta. \quad (27)$$

Since Ξ_{min}^Δ , by definition, only traces over minimally disruptive parameters, we know, for any $\delta \in (0, \delta^2(F)]$, that

$$C_\bullet(\gamma^*(\delta)) \geq C_\bullet(\Xi_{min}^\Delta(\delta)).$$

Furthermore the inequality is strict for at least one such δ , in order that γ^* itself does not only trace over minimally disruptive parameters. For (27) to hold, we therefore require

$$\int_0^{\delta^1(F)} \left\| \frac{d}{d\delta} [\gamma^*(\delta)] \right\|_2 d\delta < \int_0^{\delta^2(F)} \left\| \frac{d}{d\delta} [\Xi_{min}^\Delta(\delta)] \right\|_2 d\delta.$$

The LHS and RHS of the above inequality are equal to the lengths of the curves γ^* and Ξ_{min}^Δ , respectively. So γ^* must be strictly shorter than Ξ_{min}^Δ . However, this contradicts (25), which demands that both γ^* and Ξ_{min}^Δ have length F . So γ^* cannot exist. \square

C Necessary conditions on minimally disruptive curves

C.1 The solution of an abstract optimal control problem

We begin by providing the outline of a standard solution to an abstract optimal control problem. This problem is later specialised to provide necessary conditions on the minimally disruptive curves described in the paper. A more detailed solution of the abstract problem is provided in e.g. [7, Ch.3].

Consider a controlled dynamical system of the following form:

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t)) & x(0) &= x_0 \\ x(t) &\in \mathbb{R}^n & u(t) &\in \mathbb{R}^m \end{aligned}$$

Here $t \in \mathbb{R}^+$ denotes time, while $u(t)$ is a vector of time-varying control inputs, which the user may vary so as to affect dynamics of the state, $x(t)$. Dotted variables denote time derivatives. Given a globally nonnegative potential $L(x, u) \in \mathbb{R}^+$, we wish to find necessary conditions on $u(t)$ for minimisation of the following action:

$$J(u) = \int_0^{t_F} L(x(t), u(t)) dt, \quad (28a)$$

for some fixed time t_F , and subject to constraints of the form

$$K_1(x(t), u(t)) \leq 0, \quad K_2(x(t), u(t)) = 0. \quad (28b)$$

Note that along trajectories of the dynamical system, (28a) can be rewritten as

$$J(u) = \int_0^{t_F} L(x(t), u(t)) + \langle \lambda(t), [\dot{x}(t) - f(x(t), u(t))] \rangle dt,$$

for any function $\lambda : [0, t_F] \rightarrow \mathbb{R}^n$. Consider the Hamiltonian:

$$H(x, u, \lambda) = L(x, u) + \langle \lambda(t), f(x, u) \rangle. \quad (29)$$

Integrating by parts with respect to t gives

$$\begin{aligned} J(u) &= \int_0^{t_F} H(x(t), u(t), \lambda(t)) + \langle \dot{\lambda}(t), x(t) \rangle dt \\ &\quad + \langle \lambda(0), x_0 \rangle - \langle \lambda(t_F), x(t_F) \rangle. \end{aligned}$$

From this, the first variation of J , given fixed x_0 and free $x(t_F)$, is given by

$$\begin{aligned} \delta J(\delta u) &= \int_0^{t_F} \left\langle \left[\frac{\partial H}{\partial x}(x(t), u(t), \lambda(t)) + \dot{\lambda}(t) \right], \delta x(t) \right\rangle dt \\ &\quad + \int_0^{t_F} \left\langle \frac{\partial H}{\partial u}(x(t), u(t), \lambda(t)), \delta u(t) \right\rangle dt - \lambda(t_F) \delta x(t_F). \end{aligned}$$

Since $\lambda(t)$, which we term the costate, has thus far been arbitrary, we impose

$$\dot{\lambda}(t) = -\frac{\partial H}{\partial x}(x(t), u(t), \lambda(t)) \quad \lambda(t_F) = 0,$$

to get

$$\delta J(\delta u) = \int_0^{t_F} \left\langle \frac{\partial H}{\partial u}(x(t), u(t), \lambda(t)), \delta u(t) \right\rangle dt.$$

The form of the costate derivative also ensures time-invariance of the Hamiltonian, as can be seen by explicit differentiation.

Without the presence of the constraints (28b), the condition $\delta J = 0$ for all perturbations $\delta u : [0, t_F] \rightarrow \mathbb{R}^m$, would be a necessary condition for first order optimality. This is equivalent to the condition $\frac{\partial H}{\partial u}(x(t), u(t), \lambda(t)) = 0$ for all $t \in [0, t_F]$. As we are in fact looking for a constrained minimum,

we only need deal with perturbations δu preserving the constraints. For this, we repeat the previous arguments, but apply the Karush-Kuhn-Tucker (KKT) conditions on constrained optima at each timepoint. This gives a modified Hamiltonian:

$$\tilde{H}(x, u, \lambda) = H(x, u, \lambda) + \mu_1(t)K_1(x, u) + \mu_2(t)K_2(x, u), \quad (30)$$

where $\mu_1(t)$ is a time-varying KKT multiplier satisfying

$$\mu_1(t) \geq 0 \quad \mu_1(t)K_1(x(t), u(t)) = 0,$$

and $\mu_2(t)$ is a time varying Lagrange multiplier. The necessary conditions on optimality become:

$$\dot{\lambda}(t) = -\frac{\partial \tilde{H}}{\partial x}(x(t), u(t), \lambda(t)) \quad (31a)$$

$$\lambda(t_F) = 0 \quad (31b)$$

$$\frac{\partial \tilde{H}}{\partial u}(x(t), u(t), \lambda(t)) = 0 \quad \forall t \in [0, t_F]. \quad (31c)$$

Note that the value of $\tilde{H} - H$ is null by definition, although their partial derivatives differ. The value of \tilde{H} itself is set by (31b) and (29), which together imply that

$$\tilde{H}(x(t), u(t), \lambda(t)) = L(x(t_F), u(t_F)). \quad (32)$$

C.2 Specialisation to minimally disruptive curves

Consider a model with nominal parameter vector θ^* , and a non-negative cost function $C_\bullet(\theta)$, where C_\bullet can represent either C_{θ^*} or C_D , depending on the presence of observational data. In the paper, a minimally disruptive curve γ^* over $[0, \Delta]$ and relative to θ^* , is defined as a solution of the following optimization problem:

$$\gamma^* = \min_{\gamma} \int_0^F C_\bullet(\gamma(s)) \|\gamma'(s)\|_2 ds$$

subject to

$$\begin{aligned} \gamma(0) &= \theta^* & \langle \gamma'(s), \gamma(s) - \theta^* \rangle &> 0 \\ \gamma &\in \mathcal{C}^1 & \|\gamma'(s)\|_2 &= 1. \end{aligned}$$

This is a special case of the optimization (28). In particular, the set of system states $x \in \mathbb{R}^n$ in (28) is replaced by the set of parameter vectors $\theta \in \mathbb{R}^q$, and the time dependent state trajectory $x(t)$ is replaced by the length-scale dependent curve $\gamma(s)$. We take complete control over $\gamma'(s)$ by setting $\gamma'(s) = u$. The Lagrangian forming the integrand of (28a) is then given by $C_\bullet(\gamma(s))\|u\|_2$. The constraints K_1 and K_2 are given by

$$K_1(\gamma(s), u(s)) = -\langle u(s), \gamma(s) - \theta^* \rangle \quad (33a)$$

$$K_2(\gamma(s), u(s)) = \|u\|_2 - 1. \quad (33b)$$

and the modified Hamiltonian (30) is

$$\begin{aligned} \tilde{H}(\gamma(s), u(s), \lambda(s)) &= C_\bullet(\gamma(s)) + \langle \lambda(s), u(s) \rangle \\ &\quad - \mu_1(s) \langle u(s), \gamma(s) - \theta^* \rangle + \mu_2(s) (\|u\|_2 - 1). \end{aligned}$$

Thus the optimality condition (31c) becomes:

$$2\mu_2(s)u(s) - \mu_1(s)(\gamma(s) - \theta^*) + \lambda(s) = 0, \quad (34)$$

while (32) gives

$$\tilde{H}(\gamma(s), u(s), \lambda(s)) = C_\bullet(\gamma(F))$$

If we take the inner product of (34) with $u(s)$, while heeding the constraints, we see that

$$2\mu_2(s) = C_\bullet(\gamma(s)) - \tilde{H}(\gamma(s), u(s), \lambda(s)).$$

Given $\gamma(s)$, $\lambda(s)$ and $C_{\theta^*}(\gamma(s))$, both $\mu_1(s)$ and $u(s)$ can now found as the solution of a system of linear equations. Given a user-specified $u(0)$ and $C_{\theta^*}(\gamma(F))$, the pair of curves $\{\gamma(s), \lambda(s)\}$ can be explicitly evolved as an ordinary differential equation (ODE). Note that F need not be defined explicitly.

References

- [1] J. Anderson and A. Papachristodoulou. On validation and invalidation of biological models. *BMC bioinformatics*, 10(1):132, 2009.
- [2] M. Anguelova. *Observability and identifiability of nonlinear systems with applications in biology*. PhD thesis, Chalmers University of Technology, 2007.

- [3] J. F. Apgar, D. K. Witmer, F. M. White, and B. Tidor. Sloppy models, parameter uncertainty, and the role of experimental design. *Molecular bioSystems*, 6(10):1890–900, 2010.
- [4] R. Bellman and Åström, K.J. On structural identifiability. *Mathematical Biosciences*, 7(3):329–339, 1970.
- [5] S. Brenner. Sequences and consequences. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1537):207–212, 2010.
- [6] K. S. Brown and J. P. Sethna. Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68(2):021904, 2003.
- [7] A.E. Bryson and Ho Y-H. *Applied optimal control: optimization, estimation and control*. Taylor & Francis, 1975.
- [8] G. Casella and R.L. Berger. *Statistical inference*. Duxbury, 2nd edition, 2002.
- [9] O-T. Chis, J.R. Banga, and E. Balsa-Canto. Structural identifiability of systems biology models: a critical comparison of methods. *PloS one*, 6(11):e27755, 2011.
- [10] O-T. Chis, J.R. Banga, and E. Balsa-Canto. Sloppy models can be identifiable. *arXiv preprint*, (arXiv:1403.1417), 2014.
- [11] L. Denis-Vidal and G. Joly-Blanchard. An easy to check criterion for (un)identifiability of uncontrolled systems and its applications. *IEEE Transactions on Automatic Control*, 45(4):768–771, 2000.
- [12] L. Denis-Vidal and G. Joly-Blanchard. Equivalence and identifiability analysis of uncontrolled nonlinear dynamical systems. *Automatica*, 40(2):287–292, 2004.
- [13] N.D. Evans, M.J. Chappell, and K.R. Godfrey. Identifiability of uncontrolled nonlinear rational systems. *Automatica*, 38(10):1799–1805, 2002.
- [14] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS computational biology*, 3(10):1871–78, 2007.

- [15] S. Hengl, C. Kreutz, J. Timmer, and T. Maiwald. Data-based identifiability analysis of non-linear dynamical models. *Bioinformatics*, 23(19):2612–8, 2007.
- [16] R. Hermann and A. J. Krener. Nonlinear controllability and observability. *IEEE Transactions on automatic control*, 22(5):728–740, 1977.
- [17] K.E. Hines, T. R. Middendorf, and R. W. Aldrich. Determination of parameter identifiability in nonlinear biophysical models: A Bayesian approach. *The Journal of General Physiology*, 143(3):401–16, 2014.
- [18] M. Joshi, A. Seidel-Morgenstern, and A. Kremling. Exploiting the bootstrap method for quantifying parameter confidence intervals in dynamical systems. *Metabolic engineering*, 8(5):447–55, 2006.
- [19] H Khalil. *Nonlinear Systems Third Edition*. Prentice Hall Upper Saddle River, 3rd edition, 2002.
- [20] C. Letellier, L. A. Aguirre, and J. Maquet. Relation between observability and differential embeddings for nonlinear dynamics. *Physical Review E*, 71(6):066213, 2005.
- [21] T. Lipniacki, P. Paszek, and A.R. Brasier. Mathematical model of NF- κ B regulatory module. *Journal of theoretical biology*, pages 195–215, 2004.
- [22] Y-Y. Liu, J-J. Slotine, and A-L. Barabási. Observability of complex systems. *Proceedings of the National Academy of Sciences*, 110(7):2460–5, 2013.
- [23] L. Ljung and T. Glad. On global identifiability for arbitrary model parametrizations. *Automatica*, 30(2):265–276, 1994.
- [24] G. Margaria, E. Riccomagno, M. J. Chappell, and H. P. Wynn. Differential algebra methods for the study of the structural identifiability of rational function state-space models in the biosciences. *Mathematical Biosciences*, 174(1):1–26, 2001.
- [25] MathWorks. Matlab and statistics toolbox release 2013b, 2013.
- [26] N. Meshkat, M. Eisenberg, and J. J. Distefano. An algorithm for finding globally identifiable parameter combinations of nonlinear ODE models using Gröbner Bases. *Mathematical biosciences*, 222(2):61–72, 2009.

- [27] C. G. Moles, P. Mendes, and J. R. Banga. Parameter estimation in biochemical pathways: a comparison of global optimization methods. *Genome research*, 13(11):2467–2474, 2003.
- [28] T. P. Peixoto. Model Selection and Hypothesis Testing for Large-Scale Network Models with Overlapping Groups. *Physical Review X*, 4(1):011033, 2014.
- [29] H Pohjanpalo. System identifiability based on the power series expansion of the solution. *Mathematical biosciences*, 1978.
- [30] V. Raia, M. Schilling, M. Böhm, B. Hahn, A. Kowarsch, A. Raue, C. Sticht, S. Bohl, M. Saile, P. Möller, N. Gretz, J. Timmer, F. Theis, W.-D. Lehmann, P. Lichter, and U. Klingmüller. Dynamic mathematical modeling of IL13-induced signaling in Hodgkin and primary mediastinal B-cell lymphoma allows prediction of therapeutic targets. *Cancer research*, 71(3):693–704, 2011.
- [31] D.V. Raman, J. Anderson, and A. Papachristodoulou. On the performance of nonlinear dynamical systems under parameter perturbation. *Automatica*, 63:265–273, 2016.
- [32] A. Raue, J. Karlsson, and M.P. Saccomani. Comparison of approaches for parameter identifiability analysis of biological systems. *Bioinformatics*, 2014.
- [33] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–9, 2009.
- [34] A. Sedoglavic. A probabilistic algorithm to test local algebraic observability in polynomial time. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*. ACM, 2001.
- [35] M. K. Transtrum, B. B. Machta, and J.P. Sethna. Geometry of nonlinear least squares with applications to sloppy models and optimization. *Physical Review E*, 83(3):036701, 2011.
- [36] M. K. Transtrum and P. Qiu. Model reduction by manifold boundaries. *Physical Review Letters*, 113(9):098701, 2014.
- [37] M.K. Transtrum, B. B. Machta, and J. P. Sethna. Why are nonlinear fits to data so challenging? *Physical Review Letters*, 104(6):060201, 2010.

- [38] S. Vajda, K. R. Godfrey, and H. Rabitz. Similarity transformation approach to identifiability analysis of nonlinear compartmental models. *Mathematical Biosciences*, 93(2):217–248, 1989.
- [39] S. Vajda, H. Rabitz, E. Walter, and Y. Lecourtier. Qualitative and quantitative identifiability analysis of nonlinear chemical kinetic models. *Chemical Engineering Communications*, 83(1):191–219, 1989.
- [40] M. Vallisneri. Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Physical Review D*, 77(4):042001, 2008.
- [41] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer, V. Elser, and J. P. Sethna. The sloppy model universality class and the Vandermonde matrix. *Physical Review Letters*, 97(15):150601, 2006.
- [42] A. J. Whalen, S. N. Brennan, T. D. Sauer, and S. J. Schiff. Observability and Controllability of Nonlinear Networks: The Role of Symmetry. *Physical Review X*, 5(1):011005, 2015.
- [43] X. Xia and C. H. Moog. Identifiability of nonlinear systems with application to HIV/AIDS models. *IEEE Transactions on Automatic Control*, 48(2):330–336, 2003.
- [44] J. W.T. Yates, N.D. Evans, and M.J. Chappell. Structural identifiability analysis via symmetries of differential equations. *Automatica*, 45(11):2585–2591, 2009.

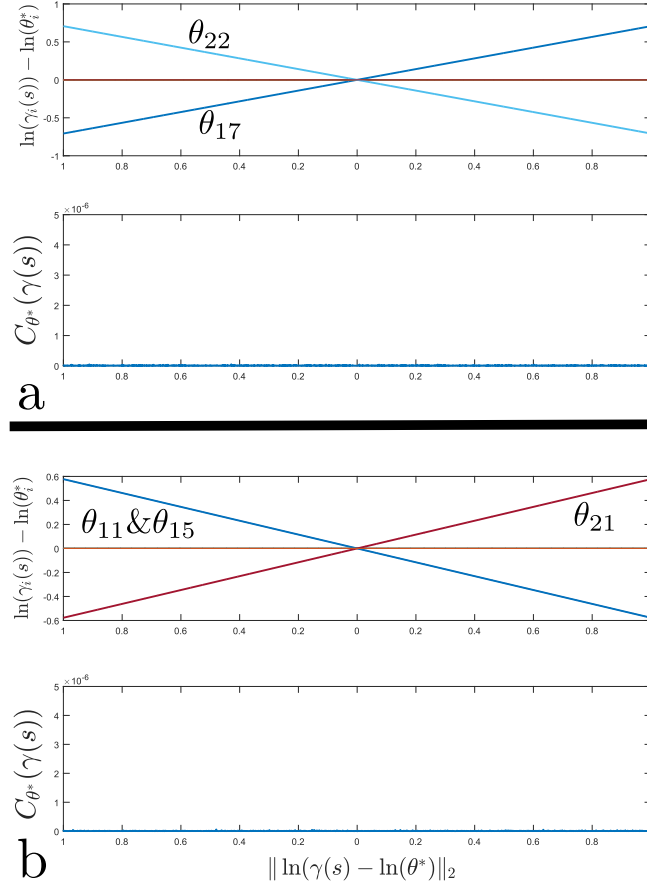


Figure 7: The x -axes of all plots in the panels depict the distance, in log-space, between the nominal parameter vector θ^* , and that of $\gamma(s)$, as s varies. Each line on the top-half figures of a and b represents the relative change of a particular (labelled) parameter component θ_i , from θ_i^* , as s varies. The bottom half figures represent the cost associated with the parameter vector $\gamma(s)$, relative to θ^* , as s varies. Note that in the top half of figure b, both θ_{11} and θ_{15} are represented by the same line: they are overlaid so closely as to be indistinguishable.

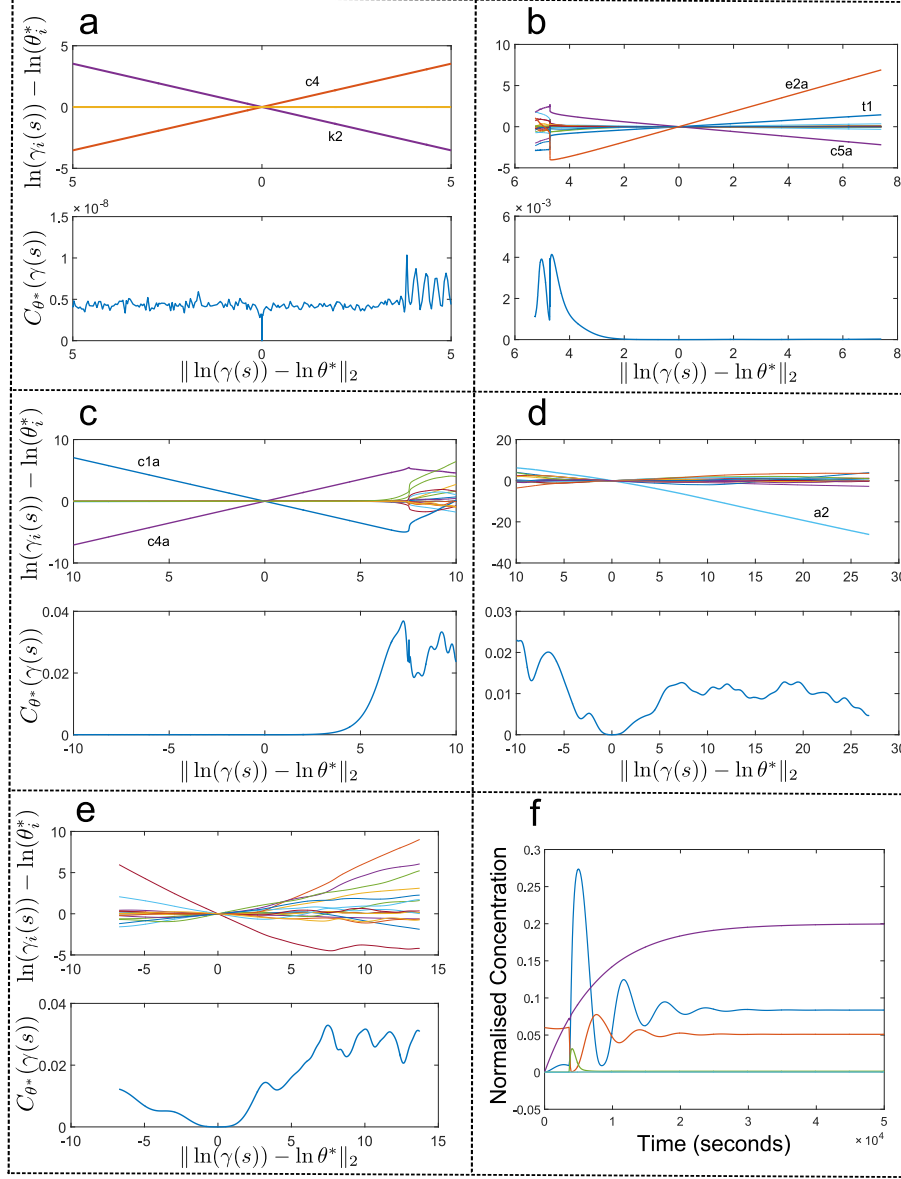


Figure 9: Panels a-e depict structural and practical unidentifiabilities in the NF- κ B model from [21]. The x -axes of all plots in these panels depict the distance, in log-space, between the nominal parameter vector θ^* , and that of $\gamma(s)$, as s varies. Each line on the top-half figures of these panels represents the relative change of a particular parameter component θ_i , from θ_i^* , as s varies. The bottom half figures represent the cost associated with the parameter vector $\gamma(s)$, relative to θ^* , as s varies. Panel f provides the normalised time-course of the observed quantities in the model.